

# Chapter 7: Overview of the Enstore Servers

In this chapter we describe the software modules that act as Enstore servers and the libraries with which they interact. The servers include:

- File Clerk (FC)
- Volume Clerk (VC)
- Library Manager (LM)
- Mover (MV)
- Media Changer (MC)
- Configuration Server (CS)

All of the above-listed servers must be running in order for data reads and writes to succeed.

The Enstore monitoring framework includes:

- Inquisitor
- Alarm Server (AS)
- Log Server (LS)
- Event Relay (ER)
- Monitor Server
- Accounting Server
- Drivestat Server

Typically, data transfer can still take place even if any of these monitoring systems is down.

## 7.1 File Clerk

---

The File Clerk (FC) is a server that tracks files in the system. It manages a database of metadata for each data file in the Enstore system. The metadata includes the file's name, its unique identifier (the bit file ID, or bfid, that the FC itself assigns to each new file), the volume on which it resides, and so on. You can get information on specific files using the **enstore file** command; see section 8.1 *enstore file*.

## 7.2 Volume Clerk

---

The Volume Clerk (VC) is a server that stores and administers storage volume (tape) information. You can get information on specific volumes using the **enstore volume** command; see section 8.4 *enstore volume*.

## 7.3 Library Manager

---

A Library in Enstore is comprised of both the physical media and a robot arm used to mount the media in attached drives. An Enstore library is typically called a robot. A library/robot interfaces to software that controls the robot arm (the Media Changer, see section 7.5). Each library can contain a variety of media types and employ different types of media drives.

A Virtual Library (VL) is a subset of an Enstore library. It can contain one and only one type of media. A Library Manager (LM) is a server which is bound to a single Virtual Library (VL), and controls what happens within that VL. We speak of bound “LM-VL pairs”. An LM receives requests for file reads and writes from the user, stores these unassigned requests in a queue, prioritizes them, and dispatches the requests to a Mover for actual data transfer to and from its VL.

There may be many LM-VL pairs in an Enstore system. There may be more than one LM-VL pair for each media type, but not vice-versa. For example, given an STK Powderhorn library holding 20, 60 and 200 GB media, Enstore would need to divide it into at least three LM-VL pairs.

You can get information on specific library managers using the **enstore library** command; see section 8.2 *enstore library*.

## 7.4 Mover

---

A Mover (MV) is a process responsible for efficient data transfer between the **encp** process and a single, assigned media drive in a library (robot). The Mover receives instructions from a Library Manager (LM) on how to satisfy the users’ requests. The Mover sends instructions to the Media Changer (MC) (described in section 9.5) that services the Mover's assigned drive in order to get the proper volume mounted.

A mover can be configured to serve multiple LMs<sup>1</sup>. Allowing flexible LM assignment has two benefits:

- First, since a virtual library (an LM-VL pair) handles only one type of media, a drive which can handle multiple types of media (e.g., different capacity media) can be shared by multiple LM-VL pairs without a static partitioning of the system.
- Secondly, suppose user groups A and B want to share the capacity of a VL, in which half the tapes belong to group A and the other half to group B. You want to guarantee that groups A and B each get one third of the tape drives, and that the last third is shared. To do this, your administrator can configure the Movers to partition resources in the VL, and assign an LM to each type of use.

## 7.5 Media Changer

---

The Media Changer (MC) mounts and dismounts the media into and out of drives according to requests from the Movers. One MC can serve multiple drives and thus multiple VLs (the image in section 7.4 *Mover* shows an MC associated with only one drive). When the drives are in the robot, the MC is the interface to the robotic software.

## 7.6 Configuration Server

---

The Configuration Server (CS) maintains and distributes the information about Enstore system configuration, such as the location and parameters of each Enstore component and/or server. At startup, each server asks the CS for the information it needs (e.g., the location of any other server with which it must communicate). New configurations can be loaded into the CS without disturbing the current running system.

## 7.7 Inquisitor

---

The Inquisitor monitors the Enstore servers, obtains information from them, and creates reports at regular intervals that can be viewed on the web under <http://hppc.fnal.gov/enstore/>. See section 7.10 *Event Relay* for

---

1. The media types governed by the LMs must be supported by the Mover's assigned drive.



an illustration of an Inquisitor task. The reports created by the Inquisitor include **Enstore Server Status** (section 9.4), **Encp History** (section 9.10), **Enstore Configuration** (section 9.11), and **Enstore Log Files** (section 9.13).

If the Inquisitor goes down, the **System-At-A-Glance** web page (described in section 9.3) indicates this by a red ball next to *Inquisitor*. In this case, data can still be transferred via **encp**, however, the information on the reports mentioned above doesn't continue to update!

## 7.8 Alarm Server

---

The Alarm Server (AS) maintains a record of alarms raised by other servers, and creates a report available online and described in section 9.12 *Enstore Active Alarms*. Since Enstore attempts error recovery whenever possible, it is expected that raised alarms will need human intervention to correct the problem. The AS compares each newly raised alarm with the previously raised ones (it raises a counter) in order to prevent raising the same alarm more than once. Alarm output can be configured to be sent in email messages, to a web page, and so on for notification.

## 7.9 Log Server

---

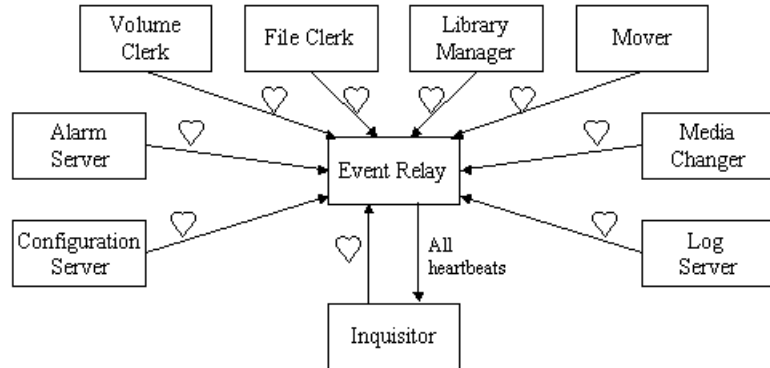
The Log Server (LS) receives messages from other processes and logs them into formatted log files available online and described in section 9.13 *Enstore Log Files*. These messages are transactional records. Log files are labeled by date. Every night at midnight, the currently opened log file gets closed and another one is opened. Logs are backed up to tape.

## 7.10 Event Relay

---

The Event Relay (ER) is a server that forwards messages based on subscription. All the Enstore servers send messages to the ER. Any server may “subscribe” to the ER in order to have messages of particular types forwarded to it.

For example, the ER periodically receives “I’m alive” messages (called *heartbeats*) from the other servers in the system. The Inquisitor (section 7.7) subscribes to the heartbeat messages, so the ER forwards these messages to it. This is illustrated in the image below:



If the ER goes down (indicated by a red ball next to *Event Relay* on the **System-At-A-Glance** web page described in section 9.3), the information on **Enstore Server Status** (see section 9.4) and the other web pages described in Chapter 9: *Monitoring Enstore on the Web* may not be accurate.

## 7.11 Monitor Server

---

The Monitor Server (MS) is available for investigating network-related problems. It attempts to mimic the communication between **encp**, the corresponding library manager, and the mover. To initiate a test of this kind, you must use the **enstore monitor** command, see section 8.3 *enstore monitor*.

## 7.12 Accounting Server

---

The accounting server is a front-end to a Postgres SQL database, and at this time is not intended for use by end users. It maintains statistical information on a running system. Such information is not essential to operations, however it can be used by admins to analyze the performance and utilization of the system for purposes of troubleshooting and future planning.

## 7.13 Drivestat Server

---

Drivestat server maintains statistical information of the drives. This is used to update the “ingest rate” plots. As with the accounting server, such information is not essential to operations, however it can be used to analyze the performance and utilization of the system for purposes of troubleshooting and future planning.